arXiv:1302.2681v1 [math.GR] 12 Feb 2013

Subgroup decomposition in $\mathsf{Out}(F_n)$: Introduction and Research Announcement

Michael Handel and Lee Mosher

February 13, 2013

Abstract

This is the introduction to a series of four papers that develop a decomposition theory for subgroups of $\mathsf{Out}(F_n)$, generalizing the theory for elements of $\mathsf{Out}(F_n)$ found in [BFH00] and [BFH05], and analogous to the decomposition theory for subgroups of mapping class groups found in [Iva92]. This work also generalizes a theorem from our earlier paper [HM09] that can be viewed as the absolute case of the main theorem of this new series. In this introduction we state the main theorems and we outline the contents of the whole series.

Given a finite type oriented surface S, Ivanov's theorem [Iva92], says that for every subgroup $\mathcal{H} < \mathcal{MCG}(S)$ of the mapping class group, either $\mathcal{H}$ has finite order, or $\mathcal{H}$ leaves invariant the isotopy class of some essential closed curve system, or $\mathcal{H}$ contains a pseudo-Anosov mapping class. Applications include the work of Bestvina and Fujiwara [BF02] that calculates bounded 2nd cohomology of arbitrary subgroups $\mathcal{MCG}(S)$ and studies homomorphisms to $\mathcal{MCG}(S)$ from lattices in higher rank Lie groups.

In this series of papers [HM13a, HM13b, HM13c, HM13d][1] we present a general study of finitely generated subgroups of $\mathsf{Out}(F_n)$. Our main result generalizes the following theorem, which was the main result of our earlier paper [HM09]:

Theorem A. *For each finitely generated*[2] *subgroup* $\mathcal{H} < \mathsf{Out}(F_n)$*, either* $\mathcal{H}$ *has a finite index subgroup that fixes the conjugacy class of some free factor of* F_n*, or* $\mathcal{H}$ *contains a fully irreducible element.*

As we shall explain, this theorem may be regarded as the "absolute case" of a broader problem on which this series of papers is focussed.

Theorem A was formulated as a strong parallel to Ivanov's Theorem, using "fully irreducible" outer automorphisms as an analogue of pseudo-Anosov mapping classes. It has found application by Bridson and Wade [BW11] in their study of homomorphisms to $\mathsf{Out}(F_n)$ of lattices in higher rank Lie groups: for any irreducible lattice in a semisimple Lie group of rank ≥ 2, any homomorphism to $\mathsf{Out}(F_n)$ has finite image.

---

[1] As of the posting of this file, Parts I and II [HM13a, HM13b] are also posted on the arXiv, while Parts III and IV [HM13c, HM13d] are still in preparation.

[2] The proof of Theorem A given in [HM09] uses that $\mathcal{H}$ is finitely generated, but the statement erred in omitting that hypothesis.

However, Theorem A has limitations. Ivanov's Theorem applies inductively: if one had already applied it to $\mathcal{H} < \mathcal{MCG}(S)$ and obtained an invariant curve system $C \subset S$ then one could restrict $\mathcal{H}$ to the components of $S - C$ and apply Ivanov's Theorem again to these restrictions. This doesn't work as well with Theorem A. Suppose we had already applied Theorem A to a subgroup $\mathcal{H} < \mathsf{Out}(F_n)$ and obtained a free factor $F < F_n$ whose conjugacy class is virtually $\mathcal{H}$-invariant. We could then apply Theorem A to "induct down into F" by restricting (a finite index subgroup of) $\mathcal{H}$ to $\mathsf{Out}(F)$, either finding a fully irreducible element or producing a still smaller free factor $F' < F$ with virtually $\mathcal{H}$-invariant conjugacy class. But that yields limited information: we cannot use Theorem A to "induct up out of F" nor to "induct in between F' and F"; in particular we will gain no information about how $\mathcal{H}$ behaves outside of F.

The main theorem (slightly simplified).

Our purpose in supplanting the preprint [HM09] with this new series is to expand the applicability of Theorem A by relativizing it. Our main result is Theorem C, which will first be stated in the slightly simplified form of Theorem C′.

Review: Free factor systems ([BFH00] Section 2.6 or see Section I.1.1.2)[3] Throughout this introduction we review material as needed from the general theory for individual outer automorphisms that is developed in works of Bestvina, Feighn, and Handel [BFH00], [FH11]. In Section I.1 of Part I [HM13a] we give an extended and thorough but terse review of this theory. We begin with free factor systems.

A *free factor system* of F_n is a finite set of the form $\mathcal{F} = \{[A_1], \ldots, [A_k]\}$ such that there is a free factorization $F_n = A_1 * \cdots * A_k * B$ where $A_1, \ldots, A_k$ are nontrivial (B may be trivial), and $[\cdot]$ denotes conjugacy class. The elements of $\mathcal{F}$ are its *components*. Inclusion of subgroups induces a partial ordering on free factor systems denoted $\mathcal{F}' \sqsubset \mathcal{F}$, which we express by saying that $\mathcal{F}'$ is *contained in* $\mathcal{F}$, or that $\mathcal{F}$ is an *extension* of $\mathcal{F}'$, or that *the pair* $\mathcal{F}' \sqsubset \mathcal{F}$ *is an extension*. Every free factor system $\mathcal{F}$ is *realized* by some marked graph G and some core subgraph $G' \subset G$, meaning that G' has no valence 1 vertices and one may list its components as $G' = C_1 \cup \ldots \cup C_k$ so that $\mathcal{F} = \{[\pi_1(C_1)], \ldots, [\pi_1(C_m)]\}$. Furthermore, any nested sequence of free factor systems can be simultaneously realized by nested core subgraphs of a single marked graph. Given a proper extension $\mathcal{F}' \sqsubset \mathcal{F}''$, if there exists a realization $G' \subset G'' \subset G$ such that the subgraph $G'' \setminus G'$ is an edge of G then $\mathcal{F}' \sqsubset \mathcal{F}''$ is a *one edge extension*, otherwise it is a *multi-edge extension*.

The group $\mathsf{Out}(F_n)$ acts naturally on free factor systems preserving the various properties described above, including the relation $\sqsubset$. For any subgroup $\Gamma < \mathsf{Out}(F_n)$ and any Γ-invariant extension of free factor systems $\mathcal{F}' \sqsubset \mathcal{F}$, we say that Γ is *reducible* relative to this extension if there exists a Γ-invariant free factor system strictly between $\mathcal{F}$ and $\mathcal{F}'$, otherwise Γ is *irreducible* relative to this extension. If furthermore each finite index subgroup of Γ is irreducible relative to the extension $\mathcal{F}' \sqsubset \mathcal{F}$ then we say that Γ is *fully irreducible* relative to this extension. These concepts apply to an individual outer automorphism $\phi \in \mathsf{Out}(F_n)$ via the cyclic subgroup $\langle\phi\rangle$, and using relative train track theory one sees that there is an attracting lamination of ϕ associated to any ϕ-invariant multi-edge extension relative to which ϕ is fully irreducible (see under *Attracting laminations* below). Also, ϕ is fully irreducible in the usual absolute sense—meaning that there are no proper, nontrivial ϕ-periodic free factor systems—if and only if $\langle\phi\rangle$ is fully irreducible with respect to the extension $\mathcal{F}' = \emptyset \sqsubset \{[F_n]\} = \mathcal{F}$.

---

[3] "Section I.X.Y.Z" or "Theorem I.V.W" refers to Section X.Y.Z or Theorem V.W of Part I [HM13a].

Here is our slightly simplified version of Theorem C:

Theorem C′. *For each finitely generated subgroup $\Gamma < \mathsf{Out}(F_n)$ and each Γ-invariant multi-edge extension of free factor systems $\mathcal{F}' \sqsubset \mathcal{F}$, if Γ is fully irreducible relative to $\mathcal{F}' \sqsubset \mathcal{F}$ then there exists $\phi \in \Gamma$ which is fully irreducible relative to $\mathcal{F}' \sqsubset \mathcal{F}$.*

The full version of Theorem C will improve on this by specifying exactly which finite index subgroup of Γ it is sufficient to consider, which allows us to remove the adjective "fully" from the hypothesis.

Rotationless versus $\mathrm{IA}_n(\mathbb{Z}/3)$ (Part II [HM13b])

Outer automorphisms of F_n can exhibit nontrivial finite permutation behavior in several ways. Given $\phi \in \mathsf{Out}(F_n)$, for example, ϕ may permute the components of a ϕ-invariant free factor system, or there may be a conjugacy class in F_n which is ϕ-periodic but not fixed by ϕ. For an individual outer automorphism one may control individual instances of this behavior by passing to a finite power.

There are two important subsets of $\mathsf{Out}(F_n)$ whose elements already have good control over finite permutations without passing to a further power: rotationless outer automorphisms ([FH11] Section 3, or see Section I.1.4); and elements of the finite index characteristic subgroup $\mathrm{IA}_n(\mathbb{Z}/3) < \mathsf{Out}(F_n)$ which, by definition, is the kernel of the natural epimorphism $\mathsf{Out}(F_n) \to H_1(F_n;\mathbb{Z}/3) \approx \mathrm{GL}(n,\mathbb{Z}/3)$. For instance, members of those two subsets satisfy the following invariance properties:

- The action of ϕ fixes each element of the finite set $\mathcal{L}(\phi)$ of attracting laminations.
- Every ϕ-periodic conjugacy class in F_n is fixed by ϕ.
- Every ϕ-periodic free factor system is fixed by ϕ and its components are fixed by ϕ.
- For any ϕ-invariant extension by free factor systems, ϕ is irreducible with respect to that extension if and only if it is fully irreducible.

For rotationless outer automorphisms, the proofs of these properties are found in [FH11]; see Fact I.1.24 for citations. Rotationless outer automorphisms have the advantage of being represented by the best relative train track maps, the so-called "completely split relative train track maps" or CTs (see [FH11] Theorem 4.28 or Theorem I.1.30). One major disadvantage is that they are not closed under the group operations.

For outer automorphisms in the subgroup $\mathrm{IA}_n(\mathbb{Z}/3)$ the proofs of these invariance properties are a significant portion of Part II [HM13b]:

Theorem B (Lemma II.4.2 and Theorems II.3.1 and II.4.1).
Each of the above invariance properties holds for each element of the subgroup $\mathrm{IA}_n(\mathbb{Z}/3)$.

The subgroup $\mathrm{IA}_n(\mathbb{Z}/3) < \mathsf{Out}(F_n)$ plays a similar role for us as the role played by the kernel of the homomorphism $\mathcal{MCG}(S) \mapsto \mathsf{Aut}(H_1(S;\mathbb{Z}/3))$ in [Iva92], namely that of providing a finite index subgroup of elements with good invariance properties. Every subgroup of $\mathsf{Out}(F_n)$ has a finite index subgroup in $\mathrm{IA}_n(\mathbb{Z}/3)$, namely its intersection; for this reason we often state our results under the restriction that the subgroup is already in $\mathrm{IA}_n(\mathbb{Z}/3)$. As an application of the Theorem B, given a subgroup $\Gamma < \mathrm{IA}_n(\mathbb{Z}/3)$, irreducibility of Γ relative to an extension of free factor systems $\mathcal{F} \sqsubset \mathcal{F}'$ is equivalent to full irreducibility relative to $\mathcal{F} \sqsubset \mathcal{F}'$.

The main theorem (full version)

Theorem C. *For each finitely generated subgroup $\Gamma < \mathrm{IA}_n(\mathbb{Z}/3)$ and each Γ-invariant multi-edge extension of free factor systems $\mathcal{F} \sqsubset \mathcal{F}'$, if Γ is irreducible relative to this extension then then there exists $\phi \in \Gamma$ which is fully irreducible relative to the extension $\mathcal{F} \sqsubset \mathcal{F}'$.*

Here is an equivalent version of Theorem C. A *filtration by free factor systems* is a strictly increasing sequence $\emptyset = \mathcal{F}_0 \sqsubset \mathcal{F}_1 \sqsubset \cdots \sqsubset \mathcal{F}_\ell = \{[F_n]\}$. Each such filtration can be realized simultaneously by some marked graph G and some nested sequence of core subgraphs $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_\ell = G$; from this one obtains the length bound $\ell \le 2n-1$. A simple induction shows for any subgroup $\Gamma < \mathsf{Out}(F_n)$, any Γ-invariant filtration by free factor systems extends to a maximal such filtration $\emptyset = \mathcal{F}_0 \sqsubset \mathcal{F}_1 \sqsubset \cdots \sqsubset \mathcal{F}_k = \{[F_n]\}$. By maximality, Γ is irreducible with respect to each extension $\mathcal{F}_{i-1} \sqsubset \mathcal{F}_i$ of this filtration. Theorem C is therefore equivalent to the following:

Theorem D. *For each finitely generated subgroup $\Gamma < \mathrm{IA}_n(\mathbb{Z}/3)$, each maximal Γ-invariant filtration by free factor systems $\emptyset = \mathcal{F}_0 \sqsubset \mathcal{F}_1 \sqsubset \cdots \sqsubset \mathcal{F}_m = \{[F_n]\}$, and each multi-edge extension $\mathcal{F}_{i-1} \sqsubset \mathcal{F}_i$, there exists $\phi \in \Gamma$ which is fully irreducible with respect to $\mathcal{F}_{i-1} \sqsubset \mathcal{F}_i$.*

Remark. We do not know whether, in Theorem D, there exists a single $\phi \in \Gamma$ which is fully irreducible with respect to each multi-edge extension $\mathcal{F}_{i-1} \sqsubset \mathcal{F}_i$, i.e. we do not know whether one can switch the order of quantification of $\mathcal{F}_{i-1} \sqsubset \mathcal{F}_i$ and ϕ. We conjecture that this can be done. The analogous statement for mapping class groups is indeed true.

Remark. We do not know whether the conclusion of Theorem C holds for all one edge extensions, although a lot can be said and this may be a relatively minor point. One edge extensions $\mathcal{F}' \sqsubset \mathcal{F}$ are of three types, which can be expressed in terms of ranks of components and in terms of realizations in a marked graph. In a *circle* extension, $\mathcal{F} = \mathcal{F}' \cup \{[B]\}$ for some rank 1 free factor B: one adds a disjoint circle to a subgraph. In a *barbell* extension, $\mathcal{F} = (\mathcal{F}' - \{[A_1],[A_2]\}) \cup \{[B]\}$ for some free factor $B < F_n$ and nontrivial free factorization $B = A_1 * A_2$ such that $[A_1],[A_2] \in \mathcal{F}'$: one attaches the endpoints of an edge to distinct components of a subgraph. In a *handle* extension, $\mathcal{F} = (\mathcal{F}' - \{[A]\}) \cup \{[B]\}$ for some free factors $A < B$ such that $\mathrm{rank}(B) = \mathrm{rank}(A) + 1$ and $[A] \in \mathcal{F}'$: one attaches the endpoints of an edge to the same component of a subgraph. If $\mathcal{F}' \sqsubset \mathcal{F}$ is a circle or barbell extension then there do not exist any free factor systems strictly between $\mathcal{F}'$ and $\mathcal{F}$, and so any element or subgroup preserving $\mathcal{F}' \sqsubset \mathcal{F}$ is fully irreducible with respect to $\mathcal{F}' \sqsubset \mathcal{F}$. However, if $\mathcal{F}' \sqsubset \mathcal{F}$ is a handle extension then there are infinitely many different free factor systems $\mathcal{F}''$ such that $\mathcal{F}' \sqsubset \mathcal{F}'' \sqsubset \mathcal{F}$; in all cases $\mathcal{F}' \sqsubset \mathcal{F}''$ is a circle extension and $\mathcal{F}'' \sqsubset \mathcal{F}$ is a barbell extension. We do not know the answer to the following question:

- If a subgroup $\Gamma < \mathrm{IA}_n(\mathbb{Z}/3)$ is irreducible with respect to a handle extension $\mathcal{F}' \sqsubset \mathcal{F}$ for which $\mathrm{rank}(A) \geq 2$, does there exist an element $\phi \in \mathcal{H}$ which is fully irreducible with respect to $\mathcal{F}' \sqsubset \mathcal{F}$?

The rest of this introduction will outline the proof of Theorem C and explain the contents of the other parts of this series, Part I [HM13a], Part II [HM13b], Part III [HM13c], and Part IV [HM13d].

The relative Kolchin theorem for $\mathsf{Out}(F_n)$ (Part II [HM13b]).

The proof of Theorem C breaks into two major steps: the relative Kolchin theorem of Part II [HM13b]; and the ping-pong arguments of Part IV [HM13d]. In [HM09], the proof of Theorem A breaks into two similar steps, the first of which is simply to cite the (absolute) Kolchin theorem from [BFH05]:

Theorem (The Kolchin type theorem for $\mathsf{Out}(F_n)$ [BFH05])**.** *For any finitely generated subgroup $\mathcal{H} < \mathsf{Out}(F_n)$, if each element of $\mathcal{H}$ is a polynomially growing outer automorphism, and if the image of each element of $\mathcal{H}$ under the natural homomorphism $\mathsf{Out}(F_n) \mapsto \mathsf{Aut}(H_1(F_n;\mathbb{Z})) \approx \mathsf{GL}(n,\mathbb{Z})$ is unipotent, then there exists an $\mathcal{H}$-invariant filtration by free factor systems $\emptyset = \mathcal{F}_0 \sqsubset \mathcal{F}_1 \sqsubset \cdots \sqsubset \mathcal{F}_k = \{[F_n]\}$ such that each extension $\mathcal{F}_{i-1} \sqsubset \mathcal{F}_n$ in this filtration is a one-edge extension.*

Recall that for ϕ to be polynomially growing means that for each automorphism $\Phi \in \mathsf{Aut}(F_n)$ representing ϕ and each $g \in F_n$, the cyclically reduced word length of $\Phi^i(g)$ is bounded above by a polynomial function of i. In topological terms this is equivalent to saying that for each marked graph G and each conjugacy class c of F_n, the length of the circuit in G representing $\phi^i(c)$ is bounded above by a polynomial. The hypothesis of the above theorem is captured by the terminology that each element $\phi \in \mathcal{H}$ is a "UPG" element of $\mathsf{Out}(F_n)$, the "U" referring to "unipotent", and the "PG" to polynomially growing. This concept is connected with $\mathrm{IA}_n(\mathbb{Z}/3)$, via the result of [BFH05] that each PG outer automorphism of F_n that is contained in $\mathrm{IA}_n(\mathbb{Z}/3)$ is UPG (see [BFH05] Proposition 3.5). The above theorem therefore applies to all subgroups $\mathcal{H} < \mathrm{IA}_n(\mathbb{Z}/3)$ consisting solely of polynomially growing outer automorphisms.

In our current, relative setting we need to generalize to subgroups $\mathcal{H} < \mathrm{IA}_n(\mathbb{Z}/3)$ such that each element of $\mathcal{H}$ leaves invariant a certain free factor system $\mathcal{F}$, allowing for the possibility that the behavior of $\mathcal{H}$ down in $\mathcal{F}$ may be very complicated, in fact elements of $\mathcal{H}$ may even have exponential growth in $\mathcal{F}$, but up outside of $\mathcal{F}$ the growth is controlled in that each element of $\mathcal{H}$ is "polynomially growing relative to $\mathcal{F}$". This defines a subset of $\mathsf{Out}(F_n)$ denoted $\mathrm{PG}^{\mathcal{F}}$. To formulate this rigorously, we review some concepts of attracting laminations.

Review: Attracting laminations ([BFH00] Section 3.1, or see Section I.1.3). Associated to each $\phi \in \mathsf{Out}(F_n)$ there is a finite, ϕ-invariant set of *attracting laminations* denoted $\mathcal{L}(\phi)$. Attracting laminations have an invariant definition which may be found in the references just cited, and they also have a definition expressed in terms of any relative train track representative of ϕ, which yields a bijection between $\mathcal{L}(\phi)$ and the EG strata of the relative train track representative (perhaps after passing to a power of ϕ that fixes each element of $\mathcal{L}(\phi)$). Attracting laminations record the exponentially growing features of ϕ, in particular, $\phi \in \mathsf{Out}(F_n)$ is polynomially growing if and only if $\mathcal{L}(\phi) = \emptyset$. Associated to each attracting lamination $\Lambda \in \mathcal{L}(\phi)$ is its *free factor support*, the smallest free factor system that "carries" each leaf of Λ, denoted $\mathcal{F}_{\text{supp}}(\Lambda)$ ([BFH00] Section 3.2, or see Section I.1.2). Distinct elements of $\mathcal{L}(\phi)$ have distinct free factor supports. For any $\phi \in \mathsf{Out}(F_n)$ there is a natural bijection $\mathcal{L}(\phi) \leftrightarrow \mathcal{L}(\phi^{-1})$ where $\Lambda^+ \leftrightarrow \Lambda^-$ if and only if $\mathcal{F}_{\text{supp}}(\Lambda^+) = \mathcal{F}_{\text{supp}}(\Lambda^-)$, in which case we say that the ordered pair $\Lambda^{\pm} = (\Lambda^+, \Lambda^-)$ is a *dual lamination pair* for ϕ. The set of dual lamination pairs for ϕ is denoted $\mathcal{L}^{\pm}(\phi)$.

If $\phi \in \mathsf{Out}(F_n)$ leaves invariant some multi-edge extension of free factor systems $\mathcal{F}' \sqsubset \mathcal{F}$ then, using standard results of relative train track theory, one easily shows the following statement: if ϕ is fully irreducible relative to that extension, then there exists an attracting lamination $\Lambda \in \mathcal{L}(\phi)$ which is supported by the free factor system $\mathcal{F}$ but not by $\mathcal{F}'$. The failure of the converse of this statement is an important point in our discussion below of how Theorem E is used in the proof of Theorem C.

Given a free factor system $\mathcal{F}$ of F_n, we let $\mathrm{PG}^{\mathcal{F}}$ denote the subset of all $\phi \in \mathsf{Out}(F_n)$ such that $\phi(\mathcal{F}) = \mathcal{F}$ and ϕ is of polynomial growth relative to $\mathcal{F}$ (see Section II.2.1). This concept may be defined in terms of laminations by saying that each $\Lambda \in \mathcal{L}(\phi)$ is supported by the free factor system $\mathcal{F}$, and it has an equivalent and more evocative definition as follows: given a marked graph G with subgraph G' realizing the free factor system $\mathcal{F}$ we have $\phi \in \mathrm{PG}^{\mathcal{F}}$ if and only if for each conjugacy class $[c]$ of F_n, letting σ_i be the circuit in G realizing $\phi^i[g]$, the number of times that σ_i crosses edges of $G \setminus G'$ is bounded above by a polynomial function of i.

The following is the main theorem of Part II [HM13b], what we call the "relative Kolchin theorem":

Theorem E (Theorem II.1.1)**.** *For any finitely generated subgroup $\mathcal{H} < \mathrm{IA}_n(\mathbb{Z}/3)$ and any free factor system $\mathcal{F}$ such that $\mathcal{H} \subset \mathrm{PG}^{\mathcal{F}}$, there exists an $\mathcal{H}$-invariant filtration by free factor systems $\mathcal{F} = \mathcal{F}_0 \sqsubset \mathcal{F}_1 \sqsubset \cdots \sqsubset \mathcal{F}_k = \{[F_n]\}$ such that for each $i = 1, \ldots, k$ the extension $\mathcal{F}_{i-1} \sqsubset \mathcal{F}_i$ is a one-edge extension.*

The proof of this theorem follows the same outline as the proof of the absolute Kolchin theorem in [BFH05]. However, many new arguments are needed, to replace arguments that work in the absolute case but not in the relative case where exponential growth is allowed as long as it is isolated in $\mathcal{F}$.

The manner in which we shall apply Theorem E is as follows. Suppose that we wish to prove Theorem C for a subgroup $\mathcal{H} < \mathrm{IA}_n(\mathbb{Z}/3)$ which is irreducible relative to an $\mathcal{H}$-invariant multi-edge extension $\mathcal{F}' \sqsubset \mathcal{F}$. We may reduce to the special case that $\mathcal{F} = \{[F_n]\}$, by restricting $\mathcal{H}$ to the individual outer automorphism groups of the components of $\mathcal{F}$. Since $\mathcal{H}$ is fully irreducible relative to the extension $\mathcal{F}' \sqsubset \mathcal{F} = \{[F_n]\}$, there exists no filtration satisfying the conclusion of Theorem E, and so $\mathcal{H} \not\subset \mathrm{PG}^{\mathcal{F}'}$. It follows that there exists $\psi \in \mathcal{H}$ that is not of polynomial growth relative to $\mathcal{F}'$, and so there exists some

$\Lambda \in \mathcal{L}(\psi)$ that is not supported by the free factor system $\mathcal{F}'$. Of course we would be done if ψ were fully irreducible relative to $\mathcal{F}' \sqsubset \mathcal{F} = \{[F_n]\}$, and while the property that Λ is not supported by $\mathcal{F}'$ may be regarded as evidence of this desired conclusion, it is not sufficient.

To work towards the conclusion of Theorem C, using ψ and Λ we shall play a game of ping-pong, although first we must build the table on which the game is to be played.

Geometric models (Part I [HM13a]).

A fully irreducible outer automorphism $\phi \in \mathsf{Out}(F_n)$ is said to be *geometric* if it is represented by a homeomorphism $h\colon S \to S$ of a compact surface S with nonempty boundary, relative to some isomorphism $\pi_1(S) \approx F_n$. It is shown in [BH92] that if this is the case then S has connected boundary and the mapping class of h is pseudo-Anosov.

In [BFH00] Definition 5.1.4, the concept of geometricity is extended from a property of a fully irreducible outer automorphism to a property of an EG (exponentially growing) stratum H_r of a relative train track map $f\colon G \to G$. The definition of geometricity is expressed in terms of the existence of what we call a "geometric model". In the very simplest case where H_r is the top stratum of G, this model is obtained by replacing H_r with a surface S, attaching all but one component of ∂S to $G \setminus H_r$, and defining a homotopy equivalence of the resulting 2-complex X that "extends" the restriction $f \bigm| (G \setminus H_r)$ over the surface S using a pseudo-Anosov homeomorphism of S.

In Part I [HM13a] we express geometricity of an EG stratum H_r in terms of a somewhat different geometric model defined by modifying Definition 5.1.4: we carefully embed H_r into S so that the entire 2-complex X deformation retracts onto G; and we modify both the way in which the components of ∂S are attached to $G \setminus H_r$ and the sense in which the homotopy equivalence of the 2-complex extends $f \bigm| (G \setminus H_r)$. As is done in [BFH00], in Part I [HM13a] we characterize geometricity of an EG stratum in various equivalent ways, one of which is expressed solely in the language of relative train track representatives.

In any relative train track map $f\colon G \to G$ representing a rotationless $\phi \in \mathsf{Out}(F_n)$ there is a bijection between the EG strata H_r of G and the set of attracting laminations of ϕ. By definition, H_r is geometric if it has a geometric model. In Proposition I.2.13 we prove that for $\Lambda \in \mathcal{L}(\phi)$ geometricity is a well-defined property of the pair (Λ, ϕ), independent of the choice of $f\colon G \to G$ and the corresponding EG stratum H_r. More precisely, if $f'\colon G' \to G'$ is another relative train track representative of ϕ with EG stratum $H_{r'}$ corresponding to Λ, then H_r is geometric if and only if $H_{r'}$ is geometric. Even more, geometricity is an invariant of the dual lamination pair $\Lambda^\pm \in \mathcal{L}^\pm(\phi)$, well-defined independent of the choice of a relative train track map representating any positive or negative power of ϕ and of the EG stratum corresponding to $\Lambda^\pm$.

As in [BFH00] and [BFH05], the distinction between geometric and nongeometric EG strata is important throughout this series of papers. In addition, we use the geometric models themselves in new ways, by applying transversality arguments in the 2-complex X. In Section I.2.2 we study self-homotopy equivalences of X with particular focus on *complementary graph of X* which is the 1-dimensional subcomplex of X identified with the closure of the complement of the interior of the surface S. In Section I.2.5 we prove Lemma I.2.15 which roughly says that any homotopy equivalence of X that preserves the complementary subgraph must also preserve the surface S up to homotopy, inducing a homeomorphism of S; this is related to Waldhausen's theorem that a homotopy equivalence of S that preserves ∂S must be homotopic to a homeomorphism of S. Lemma I.2.15 plays an important role in the proofs in Part II [HM13b] of the invariance properties for elements of $\mathrm{IA}_n(\mathbb{Z}/3)$ that are mentioned above, and hence also in the proof of Theorem E.

Vertex group systems (Part I [HM13a])

Free factor systems in F_n are of central importance in the study of $\mathsf{Out}(F_n)$. Other kinds of "subgroup systems" are also important, where by a "subgroup system" in F_n we just mean a finite set whose elements are conjugacy classes of finite rank subgroups of F_n. For example, it is proved in [GJLL98] that for any minimal action of F_n on a nondegenerate tree T with trivial arc stabilizers, the stabilizer of every point has finite rank, and there are only finitely many F_n-orbits of points whose stabilizers have positive rank; the conjugacy classes of such point stabilizers forms a subgroup system that we call the *vertex group system* of T. Proposition I.3.2, with proof suggested to us by Mark Feighn, bounds the length of any nested chain of vertex group systems.

In Part I [HM13a] we study a certain class of subgroup systems that arise naturally in connection with a geometric model X—namely the subgroup systems associated to the complementary subgraph of X and those of its subgraphs that contain ∂S—and we prove in Proposition I.3.3 that these subgroup systems are vertex group systems that are not free factor systems; this is a key component of the proof of Theorem F below. An example of this is the fact that if S is a compact surface with nonempty boundary then the components of ∂S define a vertex group system that is not a free factor system, each of whose components has rank 1. Proposition I.3.3 plays an important role in the ping-pong arguments of Part IV [HM13d].

Weak attraction theory (Part III [HM13c]).

Ping-pong arguments in topology, used for example to construct free subgroups, are usually based on a dynamical theory of attraction and repulsion applied to the individual generators of the subgroup. The ping-pong argument that is carried out in [BFH00] as part of the proof of the Tits alternative is based on the Weak Attraction Theorem, Theorem 6.0.1 of [BFH00], which describes a certain attraction–repulsion system in the space of birecurrent lines of F_n using "top" lamination pairs $\Lambda^\pm \in \mathcal{L}^\pm(\phi)$ of the generators as the attractor and repeller.

In Part III [HM13c] we develop a more expansive weak attraction theory for application to the ping-pong arguments to be carried out in Part IV [HM13d]. This theory describes attraction–repulsion systems not just on birecurrent lines but on the full space of lines $\mathcal{B} = \mathcal{B}(F_n)$, which is introduced in [BFH00] and reviewed in Section I.1.1.5. Furthermore, it allows any lamination pair $\Lambda^\pm \in \mathcal{L}(\phi)$ for any rotationless $\phi \in \mathsf{Out}(F_n)$ as the attractor–repeller pair, not just a "topmost" pair. Given any line $\ell \in \mathcal{B}$, we characterize when ℓ is attracted to Λ^+ under positive interation of ϕ, when to Λ^- under negative iteration, and when to neither. This description is broken into two theorems: Theorem F covers just the case of periodic lines; Theorem G considers the general case.

Theorem F. *For each rotationless* $\phi \in \mathsf{Out}(F_n)$ *and each dual lamination pair* $\Lambda^+ \in \mathcal{L}(\phi)$,

$\Lambda^- \in \mathcal{L}(\phi^{-1})$ *there exists a vertex group system* $\mathcal{A}_{na}(\Lambda^\pm)$*, called the* nonattracting subgroup system*, with the following properties:*

1. *For each conjugacy class* c *in* F_n *the following are equivalent:*
   - (a) c *is not weakly attracted to* Λ^+ *under iteration of* ϕ*.*
   - (b) c *is not weakly attracted to* Λ^- *under iteration of* ϕ^{-1}*.*
   - (c) c *is carried by* $\mathcal{A}_{na}(\Lambda^\pm)$*.*
2. *The lamination pair* $\Lambda^\pm$ *is geometric if and only if the vertex group system* $\mathcal{A}_{na}(\Lambda^\pm)$ *is* not *a free factor system.*

The complete description of the vertex group system is $\mathcal{A}_{\text{na}}(\Lambda^\pm)$ is given in Part III [HM13c], it is expressed very explicitly in terms of any CT $f \colon G \to G$ representing ϕ, and it generalizes the description of the subgroup system denoted $\langle Z, \rho \rangle$ in the *Weak Attraction Theorem* of [BFH00], Theorem 6.0.1. In brief, Theorem F is proved by combining relative train track theory with facts about geometric models proved in Part I [HM13a].

Theorem G, our second weak attraction result, is a general characterization of those lines $\ell \in \mathcal{B}$ that are *not* weakly attracted to Λ^+ under iteration of ϕ. We assume that both ϕ and ϕ^{-1} are rotationless, which implies that for the most important automorphisms representing either of ϕ or ϕ^{-1}—namely, the principal automorphisms—every periodic point of ∂F_n is fixed. There are three obvious sets of lines that are not weakly attracted to Λ^+: those carried by the nonattracting subgroup system $\mathcal{A}_{\text{na}}(\Lambda^\pm)$; the generic leaves of attracting laminations of ϕ^{-1}; and the "singular lines" for ϕ^{-1}, which by definition are those lines whose endpoints are nonrepelling fixed points for some single principle automorphism representing ϕ^{-1}. The statement of Theorem G refers to a binary operation of asymptotic concatenation on lines: roughly speaking, two lines which have exactly one end in common can be concatenated at that end to define another line.

Theorem G. *If* ϕ, ϕ^{-1} *are both rotationless, then a line* $\ell \in \mathcal{B}$ *is not weakly attracted to* Λ^+ *under iteration of* ϕ *if and only if* ℓ *is the asymptotic concatenation of some finite sequence of lines, each of which is in one of the three "obvious sets" described above.*

The possible concatenations which can occur in the context of this theorem are quite limited, which leads to a much more precise version of Theorem G as given in the introduction to Part III [HM13c].

Ping pong (Part IV [HM13d]).

Our earlier discussion about applying Theorem E combines with our results on weak attraction, Theorem F, in a way which motivates the ping pong strategy for proving Theorem C.

Consider a finitely generated subgroup $\mathcal{H} < \mathsf{Out}(F_n)$ and an $\mathcal{H}$-invariant multi-edge extension $\mathcal{F}' \sqsubset \mathcal{F}$ relative to which $\mathcal{H}$ is irreducible. By the earlier discussion following Theorem E, we obtain $\psi \in \mathcal{H}$ and $\Lambda_\psi \in \mathcal{L}(\psi)$ such that Λ_ψ is supported by $\mathcal{F}$ but not by $\mathcal{F}'$. To prove Theorem C it would suffice to prove that ψ is fully irreducible relative to $\mathcal{F}' \sqsubset \mathcal{F}$. But this need not be true—there are two obvious ways it could fail.

First, consider $F_{\text{supp}}(\Lambda_\psi, \mathcal{F}')$, the joint free factor support of Λ_ψ and $\mathcal{F}'$, meaning the smallest free factor system supporting Λ_ψ and containing $\mathcal{F}'$. Evidently $\mathcal{F}_{\text{supp}}(\Lambda_\psi, \mathcal{F}')$ is ψ-invariant. Evidently also we have inclusions $\mathcal{F}' \sqsubset \mathcal{F}_{\text{supp}}(\Lambda_\psi, \mathcal{F}') \sqsubset \mathcal{F}$, and the first of these inclusions is proper since Λ_ψ is not supported by $\mathcal{F}'$. But if the free factor system $\mathcal{F}_{\text{supp}}(\Lambda_\psi, \mathcal{F}')$ is "too small" in that it is properly contained in $\mathcal{F}$, then ψ is not fully irreducible relative to $\mathcal{F}' \sqsubset \mathcal{F}$.

Second, consider the nonattracting subgroup system $\mathcal{A}_{\text{na}}(\Lambda_\psi)$. Let us restrict for the moment to the case that the lamination Λ_ψ is nongeometric which, by Theorem F, is equivalent to $\mathcal{A}_{\text{na}}(\Lambda_\psi)$ being a free factor system. We may form the "meet" $\mathcal{F} \wedge \mathcal{A}_{\text{na}}(\Lambda_\psi)$ which is the unique maximal free factor system that is contained in both $\mathcal{F}$ and $\mathcal{A}_{\text{na}}(\Lambda_\psi)$ (see Fact I.1.3). Evidently $\mathcal{F} \wedge \mathcal{A}_{\text{na}}(\Lambda_\psi)$ is ψ-invariant and we have inclusions $\mathcal{F}' \sqsubset \mathcal{F} \wedge \mathcal{A}_{\text{na}}(\Lambda_\psi) \sqsubset \mathcal{F}$, the second of which is proper as one can see by an easy argument using relative train track characterization of $\mathcal{A}_{\text{na}}(\Lambda_\psi)$. But if $\mathcal{F} \wedge \mathcal{A}_{\text{na}}(\Lambda_\psi)$ is "too large", in that it properly contains $\mathcal{F}'$, then ψ is not fully irreducible relative to $\mathcal{F}' \sqsubset \mathcal{F}$. One can also carry out this analysis in the case that Λ^ψ is geometric, although the language of "meets" and "supports" needs a firmer foundation because $\mathcal{A}_{\text{na}}(\Lambda_\psi)$ is not a free factor system.

To summarize, if $\mathcal{F}_{\text{supp}}(\Lambda_\psi, F')$ is too small or if $\mathcal{A}_{\text{na}}(\Lambda_\psi)$ is too large (using the precise meanings of these terms described above) then ψ is not fully irreducible relative to $\mathcal{F}' \sqsubset \mathcal{F}$. A key fact for us is that the converse holds as well: if $\mathcal{F}_{\text{supp}}(\Lambda_\psi, F')$ is "not too small" and $\mathcal{A}_{\text{na}}(\Lambda_\psi)$ is "not too large" in the senses described precisely above, then ψ *is* fully irreducible relative to $\mathcal{F}' \sqsubset \mathcal{F}$.

Our strategy for proving Theorem C is to design a game of ping-pong which has the following effect: given $\mathcal{H}$, $\psi \in \mathcal{H}$, and Λ_ψ as above, if ψ is not fully irreducible relative to $\mathcal{F}' \sqsubset \mathcal{F}$ then one may use ping-pong arguments to drive down the value of $\mathcal{A}_{\text{na}}(\Lambda_\psi)$ and drive up the value of $\mathcal{F}_{\text{supp}}(\mathcal{F}', \Lambda_\psi)$ as needed, eventually producing an element of $\mathcal{H}$ which is indeed fully irreducible relative to $\mathcal{F}' \sqsubset \mathcal{F}$. Here are some details, where we break the ping-pong into two rounds: driving down $\mathcal{A}_{\text{na}}(\Lambda_\psi)$; and then driving up $\mathcal{F}_{\text{supp}}(\mathcal{F}', \Lambda_\psi)$.

Driving down $\mathcal{A}_{\text{na}}(\Lambda_\psi)$. If $\mathcal{A}_{\text{na}}(\Lambda_\psi)$ is too large, using that $\mathcal{H}$ is fully irreducible relative to $\mathcal{F}' \sqsubset \mathcal{F}$ we find a conjugating element $\theta \in \mathcal{H}$ such that θ does not fix $\mathcal{A}_{\text{na}}(\Lambda_\psi)$. Letting $\phi = \theta\psi\theta^{-1}$ and $\Lambda_\phi = \theta(\Lambda) \in \mathcal{L}(\phi)$, we have $\mathcal{A}_{\text{na}}(\Lambda_\phi) = \theta\mathcal{A}_{\text{na}}(\Lambda_\psi) \ne \mathcal{A}_{\text{na}}(\Lambda_\psi)$. We prove that for large values of the exponents the outer automorphism $\xi = \psi^l\phi^m$ has a lamination Λ_ξ with the property that $\mathcal{A}_{\text{na}}(\Lambda_\xi)$ is contained in both $\mathcal{A}_{\text{na}}(\Lambda_\psi)$ and $\mathcal{A}_{\text{na}}(\Lambda_\phi)$ and is therefore strictly smaller than both in some appropriate sense. By the chain condition on vertex systems mentioned earlier (Proposition I.3.2), we can carry out the induction argument needed to drive down the nonattracting subgroup system until it is not too large.

Driving up $\mathcal{F}_{\text{supp}}(\mathcal{F}', \Lambda_\psi)$. Having completed the previous round of ping pong, we may assume that we have $\psi \in \mathcal{H}$ and $\Lambda_\psi \in \mathcal{L}(\psi)$ such that Λ_ψ is supported by $\mathcal{F}$ but not by $\mathcal{F}'$, and $\mathcal{A}_{\text{na}}(\Lambda_\psi)$ is not too large, but $\mathcal{F}_{\text{supp}}(\mathcal{F}', \Lambda_\psi)$ is strictly smaller than $\mathcal{F}$. Again, using that $\mathcal{H}$ is fully irreducible relative to $\mathcal{F}' \sqsubset \mathcal{F}$, we find a conjugating element $\theta \in \mathcal{H}$ such that θ does not fix $\mathcal{F}_{\text{supp}}(\mathcal{F}', \Lambda_\psi)$. Again letting $\phi = \theta\psi\theta^{-1}$ and $\Lambda_\phi = \theta(\Lambda_\psi)$ we have $\mathcal{F}_{\text{supp}}(\mathcal{F}', \Lambda_\phi) = \theta\mathcal{F}_{\text{supp}}(\mathcal{F}', \Lambda_\psi) \ne \mathcal{F}_{\text{supp}}(\mathcal{F}', \Lambda_\psi)$. Again the ping-pong argument will produce $\xi = \psi^l\phi^m$ for large exponents and a lamination Λ_ξ supported by $\mathcal{F}$ but not by $\mathcal{F}'$, so that $\mathcal{A}_{\text{na}}(\Lambda_\xi)$ is unchanged, but $\mathcal{F}_{\text{supp}}(\mathcal{F}', \Lambda_\xi)$ has strictly increased. Again, induction eventually produces ξ and Λ_ξ so that $\mathcal{F}_{\text{supp}}(\mathcal{F}', \Lambda_\xi)$ is not too small, and we are done.

References

[BF02] M. Bestvina and K. Fujiwara, *Bounded cohomology of subgroups of mapping class groups*, Geom. Topol. **6** (2002), 69–89 (electronic).

[BFH00] M. Bestvina, M. Feighn, and M. Handel, *The Tits alternative for* $\mathrm{Out}(F_n)$*. I. Dynamics of exponentially-growing automorphisms.*, Ann. of Math. **151** (2000), no. 2, 517–623.

[BFH05] ———, *The Tits alternative for* $\mathrm{Out}(F_n)$*. II. A Kolchin type theorem*, Ann. of Math. **161** (2005), no. 1, 1–59.

[BH92] M. Bestvina and M. Handel, *Train tracks and automorphisms of free groups*, Ann. of Math. **135** (1992), 1–51.

[BW11] M. Bridson and R. Wade, *Actions of higher-rank lattices on free groups*, Compos. Math. **147** (2011), no. 5, 1573–1580.

[FH11] M. Feighn and H. Handel, *The recognition theorem for Out*(F_n), Groups Geom. Dyn. **5** (2011), 39–106, Preprint.

[GJLL98] D. Gaboriau, A. Jaeger, G. Levitt, and M. Lustig, *An index for counting fixed points of automorphisms of free groups*, Duke Math. J. **93** (1998), no. 3, 425–452.

[HM09] M. Handel and L. Mosher, *Subgroup classification in Out*(F_n), arXiv:0908:1255, 2009.

[HM13a] ———, *Subgroup decomposition in* $\mathsf{Out}(F_n)$*, Part I: Geometric models*, arXiv:1302.2378, 2013.

[HM13b] ———, *Subgroup decomposition in* $\mathsf{Out}(F_n)$*, Part II: A relative Kolchin theorem*, arXiv:1302.2379, 2013.

[HM13c] ———, *Subgroup decomposition in* $\mathsf{Out}(F_n)$*, Part III: Weak Attraction Theory*, in preparation, 2013.

[HM13d] ———, *Subgroup decomposition in* $\mathsf{Out}(F_n)$*, Part IV: Relatively irreducible subgroups*, in preparation, 2013.

[Iva92] N. V. Ivanov, *Subgroups of Teichmüller modular groups*, Translations of Mathematical Monographs, vol. 115, Amer. Math. Soc., 1992.